



CORNING

Growth in the Data Center Interconnect Segment

By David Hessong and Derek Whitehurst, Corning
As published in LIGHTWAVE® Magazine

The data center interconnect application has emerged as an important and fast-growing segment in the network landscape. This article will explore several of the reasons for this growth, including market changes, network architecture changes, and technology changes.

The huge growth of data has driven the construction of data center campuses, notably hyperscale data centers. Now, several buildings on a campus must all be connected with adequate bandwidth. How much bandwidth you may ask? To keep the information flowing between the data centers in a single campus, each data center could be transmitting to other

data centers at capacities of up to 200 Tbps today, with higher bandwidths necessary for the future.

The next question is what is driving the need for this massive amount of bandwidth between the buildings on a campus. This can be explained by increasing trends unfolding on two fronts. First, the exponential east-west traffic growth is being bolstered by machine-to-machine communication. The second trend is related to the adoption of flatter network architectures, such as leaf-and-spine or Clos networks. The goal is to have one large network fabric on the campus, thus requiring large amounts of connectivity between facilities.

Traditionally, a data center was architected on a three-tiered topology, which consisted of core routers, aggregation routers, and access switches. Although mature and widely deployed, the legacy three-tiered architecture no longer addresses the increasing workload and latency demands of hyperscale data center campus environments. In response, today's hyperscale data centers are migrating to the leaf-and-spine architecture. In leaf-and-spine, the network is divided into two stages. The spine stage is used to aggregate and route packets towards the final destination, and the leaf stage is used to connect end hosts and load balance connections across the spine.

Ideally, each leaf switch fans out to every spine switch to maximize the connectivity between servers, and consequently, the network requires high-radix spine/core switches. In many environments the large spine switches are connected to a higher-level spine switch, often referred to as a campus or aggregate ion spine, to tie all the buildings in the campus together. As a result of this flatter network architecture and the adoption of high-radix switches, we expect to see the network getting bigger, more modular, and more scalable.

DCI connectivity approaches

So what is the best and most cost-effective technology to deliver this amount of bandwidth between buildings on a data center campus? Multiple approaches have been evaluated to deliver transmission rates at this level, but the prevalent model is to transmit at lower rates over many fibers. To reach 200 Tbps using this method requires more than 3,000 fibers for each data center interconnection. When you consider the necessary fibers to connect each data center to each data center in a single campus, densities can easily surpass 10,000 fibers.

A common question is when does it make sense to use DWDM or other technologies to increase the throughput on every fiber vs. constantly increasing the number of fibers.

Currently, data center interconnect applications up to 10 km often use 1310-nm transceivers that don't match the 1550-nm transmission wavelengths of DWDM systems. So the massive interconnects are supported by using high-fiber-count cable between data centers.

The next question becomes when to replace 1310-nm transceivers with pluggable DWDM transceivers in the edge switches by adding a mux/demux unit. The answer is when or if DWDM becomes a cost-effective approach for these on campus data center interconnect links. Once this happens, the same bandwidth will be achieved by using DWDM transceivers associated with much lower fiber count cables.

To arrive at an estimate for this transition, we need to look at the DWDM transceiver price and compare with incumbent transceivers. Based on price modeling for the entire link, the current prediction is that connections based on fiber-rich 1310-nm architectures will continue to be cheaper for the foreseeable future. A PSM4 (8-fiber) alternative has proven cost-effective for applications less than 2 km, another factor driving up the fiber count.

Best practices for cable selection

Now that we have established the need for extreme-density networks, it is important to understand the best ways to build them out. These networks present new challenges in both cabling and hardware. For example, using loose tube cables and single-fiber splicing is not scalable or feasible. If you are installing a 1,728-fiber cable using a loose tube design, your splicing time will be over 100 hours assuming four minutes per splice. If you use a ribbon cable configuration, splicing time drops to less than 20 hours. Although 20 hours is still a substantial time to splice a connection, it presents a huge time savings over single-fiber cable types.

At the same time, traditional cable designs present significant challenges when installed in commonly used 2- or 4-inch ducts. New cable and ribbon designs have reached the market that have essentially doubled fiber capacity in the same cross-sectional area. These products generally fall into two design approaches: One approach uses standard matrix ribbon with more closely packable subunits, and the other approach uses standard cable designs with a central or slotted core design with loosely bonded net design ribbons that can fold on each other.

When leveraging these new ribbon cable designs, network owners need to consider the hardware and connectivity options that can adequately handle and scale with these

very high-fiber counts. It can be easy to overwhelm existing hardware, and there are several key areas to think through as you develop your complete network.

How many inside plant cables will you need to install to connect to a 1,728- to 3,456-fiber outside plant cable? If you are currently using 288-fiber ribbon cables in your inside plant environment, your hardware must be able to adequately accommodate 12 to 14 cables. Your hardware would also have to manage 288 separate ribbon splices. Using any single-fiber type cables and a single-fiber splicing method in this application is not really feasible or advisable because of massive prep times and unwieldy fiber management.

Another area that can be challenging is keeping track of fibers to ensure the correct splicing. Fibers need to be adequately labeled and sorted immediately after the cable is opened because of the magnitude of fibers that must be tracked and routed. Ensuring that ribbon stacks can be bundled together and protected during loading into hardware should be a priority to avoid damaging any fibers. In most installations, a mistake that causes redoing cable prep is manageable. In the case of extreme-density networks, a mistake can have a serious impact on project completion and could cost a week delay for just one location.

Future cabling trends

What does the future hold for extreme-density networks? The most important factor right now is whether fiber counts will stop at 3,456, or if we will see these counts go even higher. Current market trends would suggest there will be requirements for counts even beyond 5,000. To maintain infrastructure that can still scale, there will be increased pressure to downsize cable sizes. With fiber packing density already approaching its physical limits, the options to further reduce the cable diameters in a meaningful way becomes more challenging.

One approach gaining traction is to use fibers where the coating size has been reduced from a typical measurement of 250 microns down to 200 microns. The fiber core and cladding sizes remain the same so there is no change in optical performance. But this reduction in size when extended across hundreds to thousands of fibers in a cable can provide a substantial reduction in total cable cross-sectional area. This technology has already been applied in some cable designs and has been used to create micro loose tube cables that are commercially available.





Development also has focused on how to best provide data center interconnect links to locations spaced much farther apart and not collocated in the same physical campus. In a typical data center campus environment, typical data center interconnect lengths are 2 km and less. These relatively short distances enable one cable to be used to provide connectivity without any splicing points. However, with data centers also being deployed around metropolitan areas to reduce latency times, distances are increasing and can approach up to 75 km. Using an extreme-density cable design in these applications makes less financial sense because of the cost to connect the high number of fibers over a long distance. In these cases, more traditional DWDM systems will continue to be the preferred choice, running over fewer fibers at 40G and higher.

We can expect the demand for extreme-density cabling to migrate from the data center environments to the access markets as network owners prepare for the coming fiber-intensive 5G rollouts. It will continue to be a challenge in the industry to develop products that can scale effectively to reach the required fiber counts while not overwhelming existing duct and inside plant environments.

Visit www.lightwaveonline.com
for additional information

CORNING

David Hessong is manager of global data center market development and Derek Whitehurst is director of global applications marketing with Corning Optical Communications (www.corning.com/opcomm).

Corning Optical Communications LLC • 4200 Corning Place • Charlotte, NC 28216 USA
800-743-2675 • FAX: 828-325-5060 • International: +1-828-901-5000 • www.corning.com/opcomm

Corning Optical Communications reserves the right to improve, enhance, and modify the features and specifications of Corning Optical Communications products without prior notification. A complete listing of the trademarks of Corning Optical Communications is available at www.corning.com/opcomm/ trademarks. All other trademarks are the properties of their respective owners. Corning Optical Communications is ISO 9001 certified.
© 2019 Corning Optical Communications. All rights reserved. LAN-2548-AEN / September 2019